



Defence Research and
Development Canada

Recherche et développement
pour la défense Canada



A Time/Frequency Spectral Based Approach to Filtering

William J.L. Read

Defence R&D Canada – Ottawa

TECHNICAL MEMORANDUM

DRDC Ottawa TM 2005-198

December 2005

Canada

A Time/Frequency Spectral Based Approach to Filtering

William J.L. Read

Defence R&D Canada – Ottawa

Technical Memorandum

DRDC Ottawa TM 2005-198

December 2005

© Her Majesty the Queen as represented by the Minister of National Defence, 2005

© Sa majesté la reine, représentée par le ministre de la Défense nationale, 2005

Abstract

This report describes a new filtering approach suitable for demodulating narrowband signals captured in digital wideband data. The new approach is based on generating the filtered and downconverted baseband signals using time/frequency spectral data, rather than the input wideband data. For communications electronic support measures systems, this spectral data will normally be generated to facilitate signal detection and parameter estimation (i.e. frequency, bandwidth, bearing, etc.). Hence the new approach is able to take good advantage of this to reduce the overall number of computations. Through comparisons with more traditional FIR filtering techniques, it is shown that the new spectral filtering approach is computationally more efficient when multiple signals are to be simultaneously demodulated.

Résumé

Le présent rapport décrit une nouvelle méthode de filtrage qui convient à la démodulation de signaux à bande étroite saisis dans des données numériques à large bande. La nouvelle approche repose sur la production des signaux en bande de base filtrés et abaissés au moyen de données spectrales temps-fréquence, plutôt qu'au moyen de données d'entrée à large bande. Dans le cas des systèmes de mesures de soutien électronique (MSE) pour les communications, ces données spectrales sont normalement générées pour faciliter la détection des signaux et l'estimation des paramètres (comme la fréquence, la largeur de bande et le relèvement). La nouvelle approche peut donc en tirer profit pour réduire le nombre global de calculs. Par des comparaisons aux techniques de filtrage FIR traditionnelles, on montre que la nouvelle méthode de filtrage spectral est plus efficace sur le plan des calculs lorsqu'il faut démoduler plusieurs signaux simultanément.

This page intentionally left blank.

Executive summary

Spectrum monitoring and communications electronic support measures systems (ESM) require considerable flexibility and sophistication to cope with increasingly complex signal environments. One system designed to meet these needs is the Military Digital Analysis System (MiDAS). It is a wideband intercept/DF system which uses a combination of personal computer based digital signal processing algorithms and a modular system architecture based on commercial-off-the-shelf (COTS) hardware. Testing carried out to date has shown the system has the potential to meet current and future requirements, and current research and development is being carried out to achieve its full functionality.

One capability of interest is the simultaneous demodulation of narrowband signals of interest. From the hardware perspective, MiDAS is well suited for this application since a wideband receiver can be set up to search the frequency band where targets of interest are expected. The difficulty is that the wider the instantaneous frequency coverage, the greater the data rate at the output of the receiver, and the greater the likelihood of intercepting extraneous interference and noise. Sorting through all this data to determine the signal parameters of the desired signals is not a trivial task, especially in a rapidly changing signal environment.

A common approach to dealing with wideband signal data is to divide the data into blocks of N samples and then convert each block of data to the frequency domain using an FFT processor. The result is a series of spectral snapshots which change over time as the signal environment changes, i.e. a time/frequency spectrum is produced. This works well for the detection and estimation of narrowband fixed frequency signals since they occupy small portions of the spectrum at any instance in time and can be separated based on frequency, while noise is spread out across the entire spectrum weakening its disruptive influence accordingly.

Once signals have been detected, estimation of the signal parameters follows in order to determine if the signal is of interest and to acquire the parameters needed for demodulation purposes. Before demodulating the signal, the data is filtered to remove noise and any other signals which might interfere with the demodulation. One way to do this is to pass the wideband data through a bandpass filter which has been matched to the frequency and bandwidth of the targeted signal.

Although this approach works well, it may not be the most computationally efficient way given that filtered data is already available in the form of the time/frequency spectrum. In this report, an approach (called the spectral filter here) is proposed which uses the time/frequency spectral data to generate the required filtered data. The main requirements to implement this approach are: a 50% overlap when generating the time/frequency spectrum (i.e. the start of each data block used for a spectral snapshot overlaps the previous block by 50%); and the sum of two FFT windows overlapped by 50% is unity. The downside of these requirements is that it doubles the amount of processing versus the case when no overlap is used (which is how MiDAS is currently

operated), and it restricts the choices of FFT windows that can be used. However, if optimum signal detection is a goal, overlapping will be required as a matter of course [1], and the restriction on the choice of FFT windows is not a debilitating condition.

The advantage of the spectral filter is that the bulk of the processing involves the generation of the time/frequency spectrum which can be done efficiently using a fast Fourier transform (FFT). Subsequent processing to extract and filter the desired signals requires very little additional computation, hence there is relatively little difference whether one or a number of signals are extracted. By way of comparison, for normal time domain based filters the processing increases as a linear function of the number of signals and may become prohibitive if more than a few signals are to be demodulated.

For comparative purposes, two other filtering approaches were also introduced for study, namely, a single stage FIR filter and a multistage FIR filter. MiDAS currently uses a multistage FIR filter. The single stage filter has the greatest flexibility in terms of designing the filter frequency response. The multistage filter has considerably less flexibility, but is computationally more efficient than the single stage filter.

In terms of designing the filter frequency response, the spectral filter falls in between the single stage and multistage filters. It is more flexible and has a better response than the multistage filter, but is still relatively inflexible when compared to the single stage filter.

In terms of computational speed, both the single stage and multistage filters are faster when only one signal is processed. However when two signals are processed simultaneously, the spectral filter is faster than the single stage filter, and when four signals are processed simultaneously, the spectral filter is also faster than the multistage filter.

The conclusion is that for typical applications involving modern wideband ESM systems, such as MiDAS, where demodulation of four or more narrowband signals is desirable and computations are at a premium, then the spectral filtering approach proposed in this report is the preferred choice.

Read, W.J.L. 2005. A Time/Frequency Spectral Based Approach to Filtering. DRDC Ottawa TM 2005-198. Defence R&D Canada – Ottawa.

Sommaire

Les systèmes de mesures de soutien électronique (MSE) pour les communications et la surveillance du spectre ont besoin d'un perfectionnement poussé et d'une grande souplesse pour fonctionner dans des conditions de transmission des signaux de plus en plus complexes. Un système conçu pour répondre à ces besoins est le système d'analyse numérique militaire (MiDAS). Il s'agit d'un système de radiogoniométrie/d'interception à large bande qui a recours à la fois à des algorithmes de traitement numérique sur ordinateur personnel et à une architecture de système modulaire sur matériel commercial courant. Les essais menés jusqu'à maintenant montrent que le système offre la possibilité de répondre aux besoins actuels et à venir, et des travaux de recherche et développement sont menés en vue de la réalisation intégrale de cette fonctionnalité.

Une capacité souhaitable est la démodulation simultanée de signaux à bande étroite présentant un intérêt. Du point de vue du matériel, le système MiDAS convient bien à cette application, du fait qu'un récepteur large bande peut être configuré pour chercher la bande de fréquences dans laquelle on s'attend à des cibles présentant un intérêt. La difficulté, c'est que, plus la fréquence de couverture instantanée est large, plus le débit est élevé à la sortie du récepteur et plus il y a de chances d'intercepter du bruit et du brouillage étrangers. Le triage des données en vue de déterminer les paramètres des signaux utiles est une tâche non négligeable, surtout lorsque les signaux changent rapidement.

Une approche commune au traitement des données relatives aux signaux à large bande consiste à diviser les données en blocs de N échantillons, puis à convertir chaque bloc de données dans le domaine fréquence au moyen d'un processeur FFT. Il en résulte une série d'instantanés spectraux, qui changent au fur et à mesure de l'évolution des conditions des signaux, c'est-à-dire qu'un spectre temps-fréquence est produit. Cela convient bien à la détection et à l'estimation des signaux à fréquence fixe à bande étroite, du fait qu'ils occupent de faibles portions du spectre à n'importe quel moment et qu'ils peuvent être séparés en fonction de la fréquence, tandis que le bruit est étalé dans l'ensemble du spectre, ce qui réduit son influence perturbatrice.

Une fois les signaux détectés, on procède à une estimation de leurs paramètres pour déterminer s'ils présentent un intérêt et pour acquérir les paramètres nécessaires aux fins de démodulation. Avant que les signaux ne soient démodulés, on filtre les données pour en retirer le bruit et tout autre signal susceptible de perturber la démodulation. À cette fin, on peut faire passer les données à large bande dans un filtre passe-bande réglé à la fréquence et à la largeur de bande du signal ciblé.

Même si cette approche donne de bons résultats, elle n'est pas nécessairement la plus efficace sur le plan des calculs, puisque les données filtrées sont déjà disponibles sous la forme de spectre temps-fréquence. Dans le présent rapport, on propose une approche (appelée filtrage spectral) qui fait appel aux données spectrales temps-fréquence pour générer les données filtrées requises. Les principales exigences de la mise en oeuvre de

cette approche sont les suivantes : chevauchement de 50% au moment de générer le spectre temps-fréquence (c'est-à-dire que le début de chaque bloc de données utilisé pour l'obtention d'un instantané spectral chevauche de 50% le bloc précédent); et la somme de deux fenêtres FFT chevauchées de 50% est l'unité. L'inconvénient de ces exigences, c'est qu'elles doublent le volume de traitement par rapport au cas où il n'y a aucun chevauchement(c'est de cette façon que le système MiDAS fonctionne à l'heure actuelle), et elles limitent la sélection de fenêtres FFT dont on peut se servir. Toutefois, si l'objectif est la détection optimale de signaux, le chevauchement constitue une exigence tout à fait normale [citer le document de référence 1], et la restriction imposée à l'égard de la sélection des fenêtres FFT ne constitue pas une condition débilante.

L'avantage du filtre spectral, c'est que la majeure partie du traitement comporte la génération du spectre temps-fréquence, ce qui peut se faire efficacement à l'aide d'une transformée de Fourier rapide (FFT). Le traitement subséquent en vue de l'extraction et du filtrage des signaux utiles requiert très peu d'opérations additionnelles, ce qui veut dire qu'il n'y a pas tellement de différence à extraire un certain nombre de signaux ou un seul. Par comparaison, dans le cas des filtres à réponse temporelle normaux, le traitement augmente en fonction linéaire du nombre de signaux et risque de devenir prohibitif s'il faut démoduler plus que quelques signaux.

Aux fins de comparaison, deux autres approches de filtrage ont également été étudiées, à savoir un filtre FIR à un seul étage et un filtre FIR à plusieurs étages. Le système MiDAS utilise actuellement un filtre FIR à plusieurs étages. C'est le filtre à un seul étage qui offre la plus grande souplesse en ce qui concerne la conception de la réponse en fréquence du filtre. Le filtre à plusieurs étages est beaucoup moins souple, mais il est plus efficace sur le plan des calculs que le filtre à un seul étage.

Pour ce qui est de la réponse en fréquence souhaitée, le filtre spectral se situe entre le filtre à un seul étage et le filtre à plusieurs étages. Il est plus souple et a une meilleure réponse que le filtre à plusieurs étages, mais il reste relativement peu souple par rapport au filtre à un seul étage.

En ce qui concerne la vitesse de calcul, le filtre à un seul étage et le filtre à plusieurs étages sont plus rapides lorsqu'un seul signal est traité. Par contre, lorsque deux signaux sont traités simultanément, le filtre spectral est plus rapide que le filtre à un seul étage et, lorsque quatre signaux sont traités simultanément, le filtre spectral est aussi plus rapide que le filtre à plusieurs étages.

La conclusion est que, pour les applications typiques comportant des systèmes MSE à large bande modernes, comme le système MiDAS, où la démodulation de quatre signaux à bande étroite ou plus est souhaitable et où la capacité de calcul est limitée, le filtrage spectral proposé dans le présent rapport représente l'approche de choix.

Read, W.J.L. 2005. A Time/Frequency Spectral Based Approach to Filtering. DRDC Ottawa TM 2005-198. R & D pour la défense Canada – Ottawa.

Table of contents

Abstract	i
Résumé	i
Executive summary	iii
Sommaire	v
Table of contents	vii
List of figures	viii
List of tables	ix
1. Introduction	1
2. Filtering using Time/Frequency Spectral Data	3
2.1 Time/Frequency Spectral Estimation	3
2.2 Generating the Time Domain Filtered Data	4
2.3 Processing Requirements	10
3. Temporal Filtering	14
3.1 Single Stage Filter	14
3.2 Multistage Filtering	15
3.2.1 Halfband Filter	15
3.2.2 Composite Filter	19
4. Comparative Results	21
4.1 Frequency Response	21
4.2 Computational Comparisons	24
5. Conclusions	27
References	29
Annex	31

List of figures

1	Time/frequency spectrum example	4
2	Impulse noise filtering example	7
3	Filter response as a function of filter length	21
4	Filter gain at a normalized frequency of ± 0.75 as a function of filter length	22
5	Frequency response of three different filtering approaches	23
6	Computation ratio as a function of filter length	25
7	Computation ratio as a function of the oversampling factor	25
8	Parameter curves (β versus n)	27

List of tables

1	Adjacent channel filter response.	9
---	---	---

This page intentionally left blank.

1. Introduction

Spectrum monitoring and communications electronic support measures systems require considerable flexibility and sophistication to cope with increasingly complex signal environments. One system designed to meet these needs is the Military Digital Analysis System (MiDAS). This is a wideband intercept/DF system which uses a combination of personal computer based digital signal processing algorithms and a modular system architecture based on commercial-off-the-shelf (COTS) hardware. Testing carried out to date has shown the system has the potential to meet current and future requirements, and current research and development is being carried out to extend its capabilities.

One capability of interest is the simultaneous demodulation of narrowband signals of interest. From the hardware perspective, MiDAS is well suited for this application since a wideband receiver can be set up to cover large portions of the frequency band where targets of interest are expected to occur. The difficulty is that the wider the frequency coverage, the greater the data rate at the output of the receiver, and the greater the likelihood of intercepting extraneous interference and noise. Sorting through all this data to determine the parameters of the desired signals is not a trivial task, especially in a rapidly changing signal environment.

A common approach to dealing with wideband data is to divide the data into blocks of N samples and then transform each block of data samples to the frequency domain using a Fast Fourier transform (FFT). The result is a series of spectral snapshots which change over time as the signal environment changes, i.e. a time/frequency representation of the spectrum is produced. This works well for the detection and estimation of narrowband fixed frequency signals since they occupy small portions of the spectrum at any instance in time and can be separated based on frequency, while noise is spread out across the entire spectrum weakening its disruptive influence accordingly.

Once signals have been detected, estimation of the signal parameters follows in order to determine if the signal is of interest and to acquire the parameters needed for demodulation. Before demodulating the signal, the data is filtered to remove noise and any other signals which might interfere with the demodulation. One way to do this is to pass the wideband data through a bandpass filter which has been matched to the frequency and bandwidth of the targeted signal.

Although this approach works well, it may not be the most computationally efficient way to do it given that filtered data is already available in the form of the time/frequency spectrum. In this report, an approach is proposed which uses the time/frequency data to generate the required filtered data. The main requirements to implement this approach are: a 50% overlap of the data blocks when generating the time/frequency spectrum (i.e. the start of each data block used for a spectral snapshot overlaps the previous block by 50%); and the sum of two FFT windows overlapped by 50% is unity.

In the rest of this report, Section 2 introduces a new approach to filtering using the

time/frequency data and discusses the processing requirements. Section 3 introduces single stage FIR filtering and multistage FIR filtering approaches for the purposes of comparison with the new approach. This section also discusses the processing requirements of the FIR approaches. Section 4 provides comparisons between the new filtering approach and the FIR approaches, both in terms of filter frequency response and processing requirements. Finally, Section 5 provides the concluding remarks.

2. Filtering using Time/Frequency Spectral Data

2.1 Time/Frequency Spectral Estimation

A block of wideband data is represented here as

$$\mathbf{x} = [x_1, x_2, \dots, x_N]^T \quad (1)$$

where the superscript T represents transpose and x_1, x_2, \dots, x_N represent the complex sequential outputs from the wideband receiver (some pre-processing, such as frequency shifting, anti-aliasing filtering, and conversion from real to complex formatted data, may have already taken place). To produce a time/frequency spectrum, the data is rearranged into columns defined by the $M \times K$ matrix

$$\mathbf{Y} = \begin{bmatrix} w_1 x_1 & w_1 x_{k+1} & w_1 x_{2k+1} & \cdots & w_1 x_{N-M+1} \\ w_2 x_2 & w_2 x_{k+2} & w_2 x_{2k+2} & \cdots & w_2 x_{N-M+2} \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ w_M x_M & w_M x_{k+M} & w_M x_{2k+M} & \cdots & w_M x_N \end{bmatrix} \quad (2)$$

where M is the number of frequency bins or channels desired for each spectral snapshot and is assumed to be a power of two (for computational efficiency of the FFT), the weighting coefficients w_1, w_2, \dots, w_M represent the FFT window, $M - k$ represents the overlap ($k = M$ for no overlap), and the number of columns is given by

$$K = \left\lfloor \frac{N - M}{k} \right\rfloor + 1 \quad (3)$$

For the approach proposed in this report, a 50% overlap is used which corresponds to $k = M/2$ where M is even-valued. The number of columns, then becomes

$$K = \left\lfloor \frac{2N}{M} \right\rfloor - 1 \quad (4)$$

For convenience, it is assumed that N is chosen to exactly satisfy the expression

$$K = \frac{2N}{M} - 1 \quad (5)$$

where K is a positive even-valued integer. Although other values of N are permitted, the extra data, $x_{M(k+1)/2+1}, \dots, x_N$ and the last column of \mathbf{Y} (if odd-numbered), would effectively be ignored.

Once \mathbf{Y} has been created, the FFT is applied to each column to convert the column time domain data to frequency domain data. The resultant output can be represented as

$$\mathbf{S} = \begin{bmatrix} s_{11} & s_{12} & \cdots & s_{1K} \\ s_{21} & s_{22} & \cdots & s_{2K} \\ \vdots & \vdots & & \vdots \\ s_{M1} & s_{M2} & \cdots & s_{MK} \end{bmatrix} \quad (6)$$

An example of the time/frequency spectrum generated using the approach described here is shown in Figure 1. Signal data was collected using MiDAS over the 27.5 to 32.5 MHz band in a laboratory environment. Three narrowband signals are featured with the third switching on at $t = 17$ mS. The power measurements are made relative to the median noise power level as measured across the frequency band.

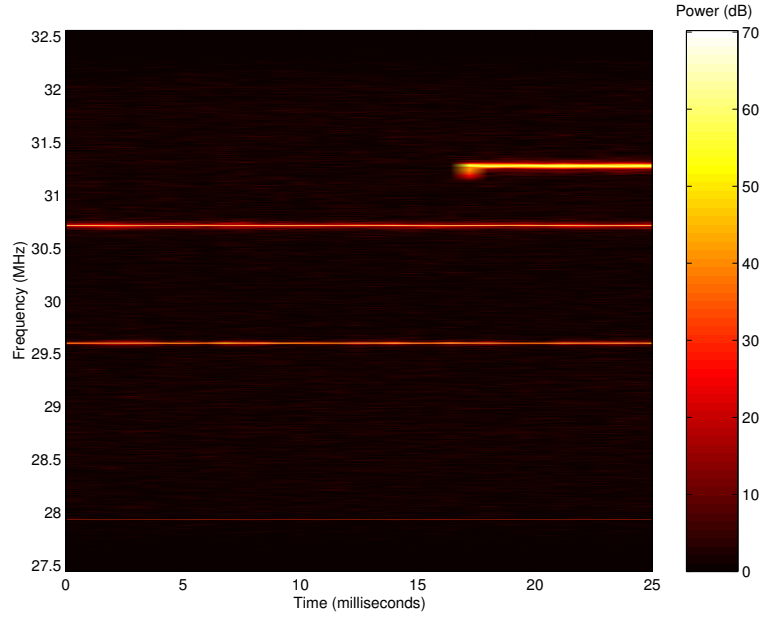


Figure 1: An example of a time/frequency spectrum produced from data featuring two narrowband signals plus a third narrowband signal switching on at $t = 17$ mS.

2.2 Generating the Time Domain Filtered Data

Given that the signal of interest has a bandwidth that is small compared to the receiver bandwidth, then filtering and decimation of the data can be implemented by limiting further processing of the data to the rows of \mathbf{S} which contain the targeted signal information. The other rows, which contain noise and interference (including other signals), are excluded. Hence, denoting the index numbers of the signal rows as $i_1, i_1 + 1, i_1 + 2, \dots, i_2$, then a filtered spectral signal matrix can be formed which is a subset of \mathbf{S} and is given by

$$\mathbf{S}_m = \begin{bmatrix} s_{i_1 1} & s_{i_1 2} & \cdots & s_{i_1 K} \\ \vdots & \vdots & & \vdots \\ s_{i_2 1} & s_{i_2 2} & \cdots & s_{i_2 K} \end{bmatrix} \quad (7)$$

where $1 < i_1 \leq i_2 < M$. An additional restriction on i_1 and i_2 is that the equation $D = M/(i_2 - i_1 + 1)$ yields an integer value for D where D is the decimation rate (the ratio of the input sampling rate to the output sampling rate of the filter).

For demodulation purposes, the time domain data is required so that \mathbf{S}_m must be

converted accordingly. This is done in two steps. The first step is to perform an inverse FFT (without using any FFT window) on the columns of \mathbf{S}_m yielding the matrix

$$\mathbf{Z} = \begin{bmatrix} z_{11} & z_{12} & \cdots & z_{1K} \\ \vdots & \vdots & & \vdots \\ z_{\frac{M}{D}1} & z_{\frac{M}{D}2} & \cdots & z_{\frac{M}{D}K} \end{bmatrix} \quad (8)$$

where $\frac{M}{D} = i_2 - i_1 + 1$. The columns of \mathbf{Z} represent filtered and decimated versions of the columns of \mathbf{Y} . In the ideal case, where the original data is noise and interference free and the spectral content of the signal-of-interest is strictly limited to rows i_1 to i_2 of the matrix \mathbf{S} , then the elements of \mathbf{Z} can be related back to the elements of \mathbf{Y} . For example, the first column of \mathbf{Y} is the inverse Fourier transform of the first column of \mathbf{S} and is given by

$$w_n x_n = \frac{1}{M} \sum_{k=1}^M s_{k1} e^{j2\pi(k-1)(n-1)/M} \quad (9)$$

for $n = 1, \dots, M$. Similarly, the first column of \mathbf{Z} is the inverse Fourier transform of the first column of \mathbf{S}_m which is given by

$$z_{m1} = \frac{D}{M} \sum_{i=1}^{\frac{M}{D}} s_{(i+i_1-1)1} e^{j2\pi(i-1)(m-1)D/M} \quad (10)$$

for $m = 1, \dots, \frac{M}{D}$. Since the signal content is limited to rows i_1 to i_2 of the matrix \mathbf{S} , then $s_{k1} = 0$ for $k < i_1$ or $k > i_2$. Given this, and replacing the index i by $k - i_1 + 1$, then

$$\begin{aligned} z_{m1} &= \frac{D}{M} \sum_{k=1}^M s_{k1} e^{j2\pi(k-i_1)D(m-1)/M} \\ &= D e^{-j2\pi(i_1-1)D(m-1)/M} \left(\frac{1}{M} \sum_{k=1}^M s_{k1} e^{j2\pi(k-1)D(m-1)/M} \right) \end{aligned} \quad (11)$$

Comparing the expression inside the large parentheses to (9) yields

$$z_{m1} = D \rho^{m-1} w_{(Dm-D+1)} x_{(Dm-D+1)} \quad (12)$$

where $\rho = e^{-j2\pi(i_1-1)D/M}$. Extending this idea to all the columns of \mathbf{Z} then, in the special noise and interference free case, \mathbf{Z} can be rewritten as

$$\mathbf{Z} = D \begin{bmatrix} w_1 x_1 & w_1 x_{(\frac{M}{2}+1)} & \cdots & w_1 x_{N-M+1} \\ w_{(D+1)} \rho x_{(D+1)} & w_{(D+1)} \rho x_{(\frac{M}{2}+D+1)} & \cdots & w_{(D+1)} \rho x_{(N-M+D+1)} \\ w_{(2D+1)} \rho^2 x_{(2D+1)} & w_{(2D+1)} \rho^2 x_{(\frac{M}{2}+2D+1)} & \cdots & w_{(2D+1)} \rho^2 x_{(N-M+2D+1)} \\ \vdots & \vdots & & \vdots \\ w_{(M-D+1)} \rho^{\frac{M}{D}} x_{(M-D+1)} & w_{(M-D+1)} \rho^{\frac{M}{D}} x_{(\frac{3M}{2}-D+1)} & \cdots & w_{(M-D+1)} \rho^{\frac{M}{D}} x_{(N-D+1)} \end{bmatrix} \quad (13)$$

The second step in deriving the time domain data is to extract the required data from \mathbf{Z} . Using (13), it is obvious that the values of $x_1, x_{D+1}, x_{2D+1}, \dots$, are easily estimated from the elements of \mathbf{Z} since ρ and the weighting coefficients, $w_1, w_{D+1}, w_{2D+1}, \dots$, are all known. The most obvious method is by directly compensating for the effects of the various multiplicative factors according to

$$x_{k_2} = \frac{z_{mn}}{Dw_{k_1} \rho^{m-1}} \quad (14)$$

where $k_1 = D(m-1) + 1$ and $k_2 = \frac{1}{2}M(n-1) + D(m-1) + 1$.

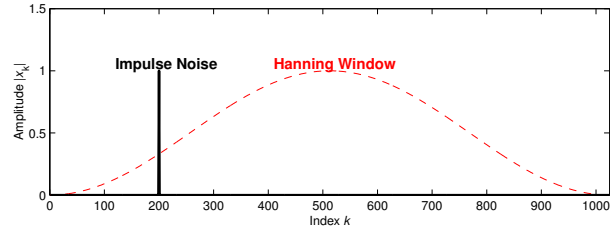
Unfortunately this simple approach has problems when, for example, effects such as impulsive noise are considered. Since the filtering process described here is additive (the result of combining several signals and noise followed by filtering is the same as individually filtering the signals and noise then combining the results), the effect of noise can be considered independently from the effect of the desired signal. Figure 2a shows the amplitude of an impulsive noise signal where the 200th value is one and all the other values are zero. The noise was converted to the frequency domain using the Fourier transform with a Hanning window. Filtering was then performed assuming the signal-of-interest occupied bins $i_1 = 65$ to $i_2 = 72$ of the spectral snapshot as shown in Figure 2b. Using the inverse Fourier transform to produce time domain data yielded the result shown in Figure 2c. The interpolated values shown in the figure were determined by performing zero padding before applying the inverse Fourier transform. Finally, Figure 2d shows the result of correcting for the Hanning window through simple division (e.g. as done in (14)). The main observation is that the noise values near the start and end of the estimated data sequence shown in Figure 2d are unrealistically large (i.e. the output data values near the start and end exceed 10000 while the largest value in the input data never exceeded 1). This is due to the fact that the corresponding window values are very small, so division using these values greatly magnifies the corresponding noise values.

The noise enhancement effect can be quantified in terms of the expected change in the overall SNR. Noting that all the estimated data values $\hat{x}_1, \hat{x}_{D+1}, \hat{x}_{2D+1}, \dots$ are affected, then the expected change in the SNR will be given by

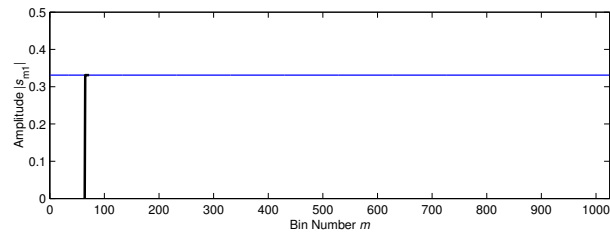
$$\Delta SNR = -10 \log \left(\frac{D}{M} \sum_{k=1}^{\frac{M}{D}} \frac{1}{w_{(kD-D+1)}^2} \right) \quad (15)$$

For most commonly used windows (e.g., Hamming, Hanning, Kaiser, Blackman, etc.), the decrease in SNR is on the order of ten's of dB. The problem can be avoided by not using any windowing (or equivalently, letting $w_k = 1$ for all values of k), however, the sidelobe response of the filtering process will be adversely affected.

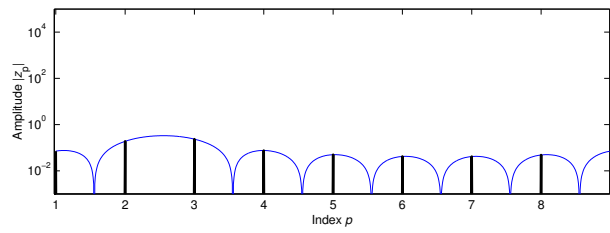
A better approach is to take advantage of the fact that each value of x_k for $k = 1, \dots, N$ appears twice in \mathbf{Y} except for the first and last $M/2$ values. That is, due to the overlapping nature of the data samples, a value of x_k appearing in the latter half of a



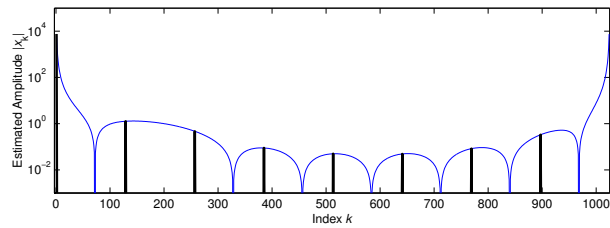
(a)



(b)



(c)



(d)

Figure 2: An example of filtering impulse noise showing (a) the input impulse noise data and the Hanning window to be applied, (b) the Fourier transform of the windowed noise (blue line) and the eight bins selected to implement the filtering (black line), (c) the inverse Fourier transform of the eight selected bins (black lines), as well as the interpolated values (blue line), (d) and the effect of correcting for the Hanning window.

column \mathbf{Y} will also appear in the first half of the next column. This overlap carries over into \mathbf{Z} , so using (13) to define an element of \mathbf{Z} as

$$z_{mn} = Dw_{k_1} \rho^{m-1} x_{k_2} \quad (16)$$

where the row index $m > M/2$, and the indices k_1 and k_2 are defined as before, then the corresponding element of \mathbf{Z} containing the same value of x_{k_2} is found in the next column and is defined as

$$z_{(m-\frac{M}{2D})(n+1)} = Dw_{(k_1-\frac{M}{2})} \rho^{(m-1-\frac{M}{2D})} x_{k_2} = (-1)^{i_1-1} Dw_{(k_1-\frac{M}{2})} \rho^{m-1} x_{k_2} \quad (17)$$

As an example, consider the case where $N \geq 128$, $M = 32$, and $D = 8$. For the first four columns of \mathbf{Z} , the elements would be defined by

$$\begin{array}{llll} +Dw_1 x_1 & & & \\ -jDw_9 x_9 & & & \\ -Dw_{17} x_{17} & +Dw_1 x_{17} & & \\ +jDw_{25} x_{25} & -jDw_9 x_{25} & & \\ & -Dw_{17} x_{33} & +Dw_1 x_{33} & \\ & +jDw_{25} x_{41} & -jDw_9 x_{41} & \\ & & -Dw_{17} x_{49} & +Dw_1 x_{49} \\ & & +jDw_{25} x_{57} & -jDw_9 x_{57} \\ & & & -Dw_{17} x_{65} \\ & & & +jDw_{25} x_{73} \end{array} \quad (18)$$

where $\rho = -j$, and the columns have been staggered here in order to line up related values of x_k .

Inspecting the two elements defined in (16) and (17), a useful quantity is given by the sum

$$z_{mn} + (-1)^{i_1-1} z_{(m-\frac{M}{2D})(n+1)} = D(w_{k_1} + w_{(k_1-\frac{M}{2})}) \rho^{m-1} x_{k_2} \quad (19)$$

In the presence of noise, this expression becomes an estimate. Rearranging, the corresponding estimate for x_k is given by

$$\hat{x}_{k_2} = \frac{z_{mn} + (-1)^{i_1-1} z_{(m-\frac{M}{2D})(n+1)}}{D(w_{k_1} + w_{(k_1-\frac{M}{2})}) \rho^{m-1}} \quad (20)$$

which is similar in form to (14), but the problem term w_{k_1} is replaced by $w_{k_1} + w_{(k_1-\frac{M}{2})}$. For most commonly used windows, the variation in $w_{k_1} + w_{(k_1-\frac{M}{2})}$ between the minimum and maximum values is less than a factor of two, considerably less than the several orders of magnitude variation in the value of w_{k_1} alone. The result is a substantial decrease in the noise enhancement effect.

Some examples of the effect on SNR for different windows are listed in Table 1. In this case the change in SNR was calculated by appropriately modifying (15) to get

$$\Delta SNR = -10 \log \left(\frac{2D}{M} \sum_{k=1}^{\frac{M}{2D}} \frac{1}{(w_{(Dk-D+1)} + w_{(Dk-D+1+\frac{M}{2})})^2} \right) \quad (21)$$

Table 1: Adjacent channel filter response.

Window	Change in SNR (dB)
Triangular	0
Hamming	0
Hanning	0
Kaiser $\beta = 5$	-0.1
Kaiser $\beta = 10$	-2.7
Kaiser $\beta = 15$	-6.2
Gaussian $\alpha = 2.5$	-0.6
Blackman	-1.8
Chebyshev $R = -70$ dB	-1.8
Bohman	-2.2
Nuttall	-3.8
Blackman-Harris	-4.0

Of the commonly used windows, the triangular window satisfies (24) exactly and the Hanning and Hamming windows nearly satisfy it (the variation between the minimum and maximum values of $w_{k_1} + w_{(k_1 - \frac{M}{2})}$ is less than 0.2%). Adjusting the weights for the other windows to satisfy (24) can be accomplished using

$$w'_{k_1} = w_{k_1} - \frac{w_{k_1} + w_{(k_1 - \frac{M}{2})} - 1}{2} \quad (22)$$

$$w'_{(k_1 - \frac{M}{2})} = w_{(k_1 - \frac{M}{2})} - \frac{w_{k_1} + w_{(k_1 - \frac{M}{2})} - 1}{2} \quad (23)$$

which adjusts both weights equally. This adjustment modifies the properties of the window, however, and tends to make the window properties more “hanning”-like.

In the event that the weights satisfy the constraint

$$w_{k_1} + w_{(k_1 - \frac{M}{2})} = 1 \quad (24)$$

the noise enhancement effect disappears completely and the estimation of x_k simplifies to

$$\hat{x}_{k_2} = \frac{z_{mn} + (-1)^{i_1-1} z_{(m - \frac{M}{2D})(n+1)}}{D\rho^{m-1}} \quad (25)$$

The filtering approach introduced here is called the *spectral filter* throughout the rest of this report. A few additional comments about using this approach are also in order. As the spectral filter has been presented up to this point, the decimation rate is tied to the desired filter bandwidth, F_{bw} , according to

$$D = \frac{M}{i_2 - i_1 + 1} = \frac{F_s}{F_{bw}} \quad (26)$$

for $i_2 > i_1$. The output data sampling rate is then given by

$$F_o = \frac{F_s}{D} \quad (27)$$

which is the Nyquist rate (i.e. $F_o = F_{bw}$). To increase the output sampling rate, while maintaining the same filter bandwidth, rows of zeros (zero padding) can be added to \mathbf{Z} at the beginning, end, or both. The resultant new decimation rate will be given by

$$D = \frac{M}{n_z + i_2 - i_1 + 1} \leq \frac{F_s}{F_{bw}} \quad (28)$$

where n_z is the number of rows of zeros added and it is assumed that n_z , i_1 , and i_2 are chosen to yield an integer value for D . The new output sampling rate is still calculated using (27).

2.3 Processing Requirements

In this section, the computational requirements of the proposed spectral filter are considered: specifically, the arithmetic requirements. Additional processing due to housekeeping, memory transfers, etc., are not considered as this becomes very dependent on the processor performing the filtering, the type of memory used, and the way the algorithm is implemented in software, which is beyond the scope of this report. For the purposes of comparing the proposed approach with other approaches discussed later in this report, it can be assumed that a generic processor is used in all cases (i.e. not optimized in a way which favours one approach over another), critical memory transfers involve cache memory only, and that all approaches are efficiently coded in software.

In order to consider the processing requirements of the proposed filtering method, it is useful to define a generic operation which can be used to quantify the computations required to carry out a given process. To do this, consider the following summation series

$$\sum_{q=1}^Q u_q v_q \quad (29)$$

which is representative of most of the mathematical computations performed in the filtering process. There are three cases to consider: both u_q and v_q are real-valued; one is real-valued and the other is complex valued; and both are complex-valued.

In the first case, where both are real-valued, Q real multiplications and $Q - 1$ additions are required to carry out (29). Given this is the simplest case, it is useful to define an operation as a single *real* multiplication plus a single *real* addition. Hence, assuming that $Q \gg 1$, which will normally be the case, then Q operations are required to carry out (29).

In the second case, where, for example u_q , is real-valued and v_q is complex valued, then (29) can be represented as

$$\sum_{q=1}^Q u_q (\Re\{v_q\} + j\Im\{v_q\}) = \sum_{q=1}^Q u_q \Re\{v_q\} + j \sum_{q=1}^Q u_q \Im\{v_q\} \quad (30)$$

which requires twice the number of operations as the first case, or $2Q$ operations.

In the third case, where both u_q and v_q are complex valued, then (29) can be represented as

$$\begin{aligned}
& \sum_{q=1}^Q (\Re\{u_q\} + j\Im\{u_q\})(\Re\{v_q\} + j\Im\{v_q\}) \\
&= \sum_{q=1}^Q \Re\{u_q\}\Re\{v_q\} - \sum_{q=1}^Q \Im\{u_q\}\Im\{v_q\} \\
&\quad + j \sum_{q=1}^Q \Re\{u_q\}\Im\{v_q\} + j \sum_{q=1}^Q \Im\{u_q\}\Re\{v_q\}
\end{aligned} \tag{31}$$

which requires four times the number of operations as the first case, or $4Q$ operations.

Finally, some of the filter processing may also include Q real multiplications without any additions. In this case the number of operations is considered to be $\frac{1}{2}Q$ for this report. This is based on the example of the Intel P4 processor which can carry out floating point additions with a throughput of 2 clock cycles and a latency of 4 clock cycles, and floating point multiplication with the same throughput but a latency of 6 clock cycles. Under ideal conditions, for a large number of pipelined instructions, the average processing time per addition will then be the same as the average processing time per multiplication. In real systems this ideal will likely not be achieved, particularly when the performance of alternate hardware implementations and different software libraries is taken into account. Hence the assumed equivalence between addition and multiplication times is only an estimate, so that the "number of computations" results that follow should also be considered estimates. However, in terms of assessing the relative performance of different filtering approaches in terms of the number of computations, the results remain valid.

Having now examined and quantified an "operation", the processing requirements of the proposed spectral filtering method can be determined. The processing can be broken into four separate steps which can be examined separately to determine the individual requirements. These four steps are:

1. the application of the window weighting coefficients in (2),
2. the application of the FFT on the columns of \mathbf{Y} to get \mathbf{S} shown in (6),
3. the inverse FFT on the columns of \mathbf{S}_m to get \mathbf{Z} shown in (8), and
4. the generation of the filtered output values from \mathbf{Z} using (14).

In the first step, the application of the weighting coefficients is a straightforward process involving the multiplication of each of $M \times K$ complex data values by a real

weighting coefficient. Given that a multiplication involving a complex and a real value is equivalent to two real multiplications, and remembering the $\frac{1}{2}$ factor used for multiplication without addition, then

$$C_1 = \frac{2MK}{2} = 2N - M \quad (32)$$

where (5) was used to replace K .

For the second step, the computational requirements of an M -point FFT are given by

$$M \log M \quad (33)$$

complex multiplies and adds [2], which is equivalent to

$$4M \log 4M \quad (34)$$

real multiplies and adds. Hence, the number of operations for converting the K columns of \mathbf{Y} to \mathbf{S} using the FFT is given by

$$C_2 = 4KM \log_2 M = (8N - 4M) \log_2 M \quad (35)$$

where (5) was used to replace K .

The third step involves an inverse FFT carried out on a subset of the rows of \mathbf{S} , namely \mathbf{S}_m . The processing requirements are given by

$$C_3 = 4K \frac{M}{D} \log_2 \frac{M}{D} = \frac{8N - 4M}{D} (\log_2 M - \log_2 D) \quad (36)$$

where the inverse FFT size is M/D and the relationship expressed in (5) was used to replace K (which represents the number of columns of \mathbf{S}_m).

For the fourth step, the number of filtered output values L which are generated is given by the number of unique values of x_k in \mathbf{Z} , or

$$L = \frac{N}{D} \quad (37)$$

Assuming that processed data from the previous and next blocks of \mathbf{Z} are available, then (20) can be used to process all the data (i.e. the data in the first half of the first column of \mathbf{Z} is combined with data in the last column of the previous version of \mathbf{Z} and the data in the last half of the last column of \mathbf{Z} is combined with data in the first column of the next version of \mathbf{Z} using an appropriately modified version of (20) without resorting to (14)). Noting that (20) involves a complex addition and multiplication (where the inverse of the denominator term can be computed beforehand), then the processing requirements are

$$C_4 = 4L = \frac{4N}{D} \quad (38)$$

In terms of the overall processing, the first two processing steps are only carried out once, regardless of the number of signals to be demodulated. Since these two steps are also being carried out for detection purposes as well, then arguably these steps could be considered as “free” processing and not included as part of the filter processing calculations. However, since current realtime systems, which generate time/frequency spectral data for detection purposes, would likely use contiguous (not overlapping) data blocks then a 50% overlap would result in a doubling of the processing ¹. Hence, the added processing attributable to the filtering process will be $\frac{1}{2}(C_1 + C_2)$.

The last two processing steps must be carried out for each signal to be demodulated. Consequently, the added processing requirements due to these two steps will be $n(C_3 + C_4)$ where n is the number of signals to be demodulated.

From this, the total number of computations required for the proposed filtering method is given by

$$\begin{aligned}
C &= \frac{1}{2}(C_1 + C_2) + n(C_3 + C_4) \\
&= \frac{1}{2}(2N - M + (8N - 4M) \log_2 M) \\
&\quad + n \left(\frac{8N - 4M}{D} (\log_2 M - \log_2 D) + \frac{4N}{D} \right) \\
&= \frac{1}{2}(2N - M)(1 + 4 \log_2 M) + \frac{4n}{D} ((2N - M)(\log_2 M - \log_2 D) + N)
\end{aligned} \tag{39}$$

For most filtering applications where signal demodulation is the objective, large blocks of input data will be used so that it is reasonable to assume $N \gg M$. This leads to a simplification in the expression for the total number of computations according to

$$\boxed{C \approx N + 4N \log_2 M + \frac{4nN}{D} (2 \log_2 M - 2 \log_2 D + 1)} \tag{40}$$

The cost of computing the Fourier coefficients $e^{j2\pi k/M}$ and $e^{-j2\pi k/M}$ for $k = 0, \dots, M - 1$ has not been considered since they only need to be computed or loaded from a table once and are used many times (i.e. since it is assumed that $N \gg M$ then $K \gg 1$) so that the additional processing load will be negligible.

¹Note that overlapping leads to improved detection performance depending on the FFT window used [1].

3. Temporal Filtering

In the following, fast variants of a single stage filter and a multiple stage filter are presented which are appropriate for demodulation purposes. Their processing requirements are also discussed in a way which makes comparisons with the spectral filter discussed in Section 2 straightforward. As discussed previously in Section 2.3, the processing requirements are restricted to arithmetic operations only.

3.1 Single Stage Filter

The basic filter operation can be represented by

$$z_i = \sum_{k=1}^M h_k y_{k+i-1} \quad (41)$$

where z_1, \dots, z_{N-M+1} represents the filtered output data, h_1, \dots, h_M represents the low pass filter coefficients (taps) which are chosen according to the desired filter response, and y_1, \dots, y_N represents the input data after the signal-of-interest has been downconverted to baseband according to

$$y_k = x_k e^{-j2\pi f_c(k-1)/F_s} \quad (42)$$

The downconversion is done prior to filtering so that only a single set of filter coefficients is required, regardless of the center frequency, f_c , of the signal. This simplifies storage and processing requirements as well as ensuring that the filter response remains constant.

In terms of the processing requirements, the downconversion features N complex multiplications, which is equivalent to $4N$ real multiplications and $2N$ real additions, or

$$C_1 = 3N \quad (43)$$

operations. It also involves the generation of the mixing term represented by $e^{-j2\pi f_c(k-1)/F_s}$. For the purposes of speed it is assumed that the appropriate values of the mixing terms are tabulated in cache memory². The number of operations will then be given by

$$C_2 = \frac{N}{2} \quad (45)$$

²This can be done by saving the sequence

$$s(n) = e^{-j2\pi n \Delta f / F_s} \quad (44)$$

for $n = 0, \dots, \frac{F_s}{\Delta f} - 1$. The quantity Δf is the frequency quantization increment and is chosen small enough that replacing f_c by $m\Delta f$ does not significantly degrade the accuracy of the filtering process. The mixing sequence is then given by $s(0), s(m), s(2m), \dots, s(mN - m)$ where the fact that $s(F_s/\Delta f + k) = s(k)$ is also used. The main computational requirement is the generation of the table index terms $0, m, 2m, \dots$ which require either a single real multiplication or single real addition per term

In terms of producing the filtered output, for many applications it is not necessary to generate every possible output value of z_k , but rather the decimated sequence represented by $z_1, z_{1+D}, z_{1+2D}, \dots$. Hence, calculating only the required output values leads to the modified filter equation given by

$$z_{1+jD} = \sum_{k=1}^M h_k y_{k+jD} \quad (46)$$

where the index $j = 0, \dots, L-1$, and z_{1+jD} is the decimated output. In this case, L output values are produced so that the number of operations required to generate the decimated sequence is

$$C_3 = 2LM \quad (47)$$

where h_1, \dots, h_M are assumed to be real-valued.

The total number of operations is simply the sum of (43), (45) and (47) times the number of signals, n , to be demodulated, or

$$C = n(C_1 + C_2 + C_3) = \frac{7nN}{2} + 2nLM \quad (48)$$

An alternate form of this expression can be derived by considering that the number of output data values will be given by

$$L = 1 + \left\lfloor \frac{N-M}{D} \right\rfloor \quad (49)$$

To simplify later comparisons between different filtering approaches, it is assumed that N is chosen so that

$$L = 1 + \frac{N-M}{D} \quad (50)$$

The total number of operations can then be rewritten as

$$C = \frac{7nN}{2} + 2nM + \frac{2nM(N-M)}{D} \quad (51)$$

Finally, using the assumption that $N \gg M$, this simplifies to

$$C \approx \frac{7nN}{2} + \frac{2nMN}{D} \quad (52)$$

3.2 Multistage Filtering

3.2.1 Halfband Filter

Faster filter implementations can be achieved by using a multistage filter. In particular, a simple halfband filter can be used to successively halve the bandwidth of the data at each stage until the desired bandwidth is achieved. The filter equation for each stage is given by modifying (46) to get

$$z_{1+jD_k}(k) = \sum_{m=1}^{M_2} h_m y_{m+jD_k}(k) \quad (53)$$

for $k = 1, \dots, K$ where K is the number of filter stages, $D_k = 2^k$ is the cumulative decimation of the data carried out by the k^{th} stage (i.e. using a halfband filters, each stage decimates the data by a factor of two so that the cumulative decimation by the k^{th} stage is 2^k), M_2 is the number of filter coefficient for each halfband filter, and $y_i(k)$ and $z_i(k)$ represent the input and output data of the stage filter, respectively (hence for $k > 1$, then $y_i(k) = z_i(k-1)$).

Before considering the computational requirements, it is first useful to determine the number of decimated output values available at the output of each stage. Letting L_1, L_2, \dots, L_K represent these quantities, then (50) can be rewritten as

$$L_k = 1 - \frac{1}{2}M_2 + \frac{1}{2}L_{k-1} \quad (54)$$

for $k = 1, \dots, K$, where L_k represents the number of filter output values for the k^{th} stage (as well as the number of input values for the next stage), and $L_0 = N$ and $L_K = L$. Beginning with L_1 , and successively solving then

$$L_1 = 1 - \frac{1}{2}M_2 + \frac{1}{2}N \quad (55)$$

$$\begin{aligned} L_2 &= 1 - \frac{1}{2}M_2 + \frac{1}{2}L_1 \\ &= \frac{3}{2} - \frac{3}{4}M_2 + \frac{1}{4}N \end{aligned} \quad (56)$$

$$\begin{aligned} L_3 &= 1 - \frac{1}{2}M_2 + \frac{1}{2}L_2 \\ &= \frac{7}{4} - \frac{7}{8}M_2 + \frac{1}{8}N \end{aligned} \quad (57)$$

and so on. For the k^{th} filter stage, the general solution is given by

$$\begin{aligned} L_k &= 1 - \frac{1}{2}M_2 + \frac{1}{2}L_{k-1} \\ &= \frac{2^{k+1} - 2}{2^k} - \frac{2^k - 1}{2^k}M_2 + \frac{1}{2^k}N \\ &= (2 - M_2) + \frac{N + M_2 - 2}{2^k} \end{aligned} \quad (58)$$

for $k = 1, \dots, K$.

The number of computations is found by taking the requirements for downconverting the signal given by (43) and adding this to the filter requirements found by applying (47) to each stage and summing the results. This is then multiplied by the number of signals, n , to be demodulated yielding

$$\begin{aligned}
C &= \frac{7nN}{2} + n \sum_{k=1}^K 2L_k M_2 \\
&= \frac{7nN}{2} + 2nM_2 \sum_{k=1}^K \left(2 - M_2 + \frac{N + M_2 - 2}{2^k} \right) \\
&= \frac{7nN}{2} + 2nM_2 K(2 - M_2) + 2nM_2(N + M_2 - 2) \sum_{k=1}^K \frac{1}{2^k} \\
&= \frac{7nN}{2} + 2nM_2 K(2 - M_2) + 2nM_2(N + M_2 - 2) \left(1 - \frac{1}{D} \right) \quad (59)
\end{aligned}$$

where

$$D = 2^K \quad (60)$$

is the decimation rate for all K stages.

It is relatively easy to show that the multistage filter implementation as discussed so far actually leads to a slower filter implementation than the single stage filter discussed previously. The main advantage of using a multistage filter implementation comes from taking advantage of the fact that almost half of the coefficients of a halfband lowpass FIR filter are zero if M_2 is odd. For example, an eleven-tap filter designed using a Hamming window has the coefficients

$$[h_1, \dots, h_{11}]^T = \begin{bmatrix} 0.00506 \\ 0 \\ -0.0419 \\ 0 \\ 0.288 \\ 0.497 \\ 0.288 \\ 0 \\ -0.0419 \\ 0 \\ 0.00506 \end{bmatrix} \quad (61)$$

where every second coefficient is zero with the exception of the central coefficient. From this it may be concluded that for $M_2 = 3, 7, 11, 15, \dots$, the number of nonzero coefficients is given by

$$M'_2 = \frac{M_2 + 3}{2} \quad (62)$$

It is not necessary to consider the cases $M_2 = 5, 9, 13, 17, \dots$ since this results in $h_1 = 0$ and $h_{M_2} = 0$ making the filter equivalent to one with $M_2 - 2$ coefficients.

The zero-value coefficients mean that fewer multiplications are required at each stage³ so (59) is modified to become

$$\begin{aligned}
C &= \frac{7nN}{2} + n \sum_{k=1}^K 2L_k M_2 \\
&= \frac{7nN}{2} + 2nM_2' K(2 - M_2) + 2nM_2'(N + M_2 - 2) \left(1 - \frac{1}{D}\right) \\
&= \frac{7nN}{2} + 2n \left(\frac{M_2 + 3}{2}\right) K(2 - M_2) + 2n \left(\frac{M_2 + 3}{2}\right) (N + M_2 - 2) \left(1 - \frac{1}{D}\right) \\
&= \frac{7nN}{2} - nK(M_2 + 3)(M_2 - 2) + n(M_2 + 3)(N + M_2 - 2) \left(1 - \frac{1}{D}\right) \quad (63)
\end{aligned}$$

where it is not necessary to modify the expression for L_k since it is based on the total number of coefficients, not the number of nonzero coefficients.

Since $N \gg M_2$ for most applications, then the number of operations simplifies to

$$\boxed{C \approx \frac{7nN}{2} + nN(M_2 + 3) \left(1 - \frac{1}{D}\right)} \quad (64)$$

For comparative purposes, the number of coefficients for each halfband filter can be related back to the number of coefficients for an equivalent single stage filter according to

$$M_2 = \frac{M - 1}{D - 1} + 1 \quad (65)$$

which follows from (58) since $N = M$ when $L = 1$. It is assumed, in this case, that M and D are chosen so that this expression yields a value for M_2 from the set $\{3, 7, 11, 15, \dots\}$. Finally, substituting this expression back into (64), yields

$$\boxed{C \approx \frac{15nN}{2} + nN \left(\frac{M - 5}{D}\right)} \quad (66)$$

which allows direct comparisons between the multistage filter and the other approaches previously discussed.

³Provided the software implementation takes advantage of the zero values – some popular software libraries do not

Note that further reductions in the number of computations are possible by modifying the number of filter taps on a stage by stage basis, where a lower number of taps is used in the earlier stages. This modified halfband filter approach, however, requires more careful design and is not considered in the rest of this report.

3.2.2 Composite Filter

One of the drawbacks of the multistage halfband filter for demodulation is that the filter bandwidth is tied to the decimation rate according to

$$F_{bw} = F_s/D \quad (67)$$

Hence, if the signal is filtered to the correct bandwidth, the output sampling rate of the filter will equal the Nyquist rate. Many digital demodulation techniques require the signal to be oversampled by a factor of two or more, so this presents a problem.

One possible solution is to use fewer filter stages to achieve the desired output sampling rate. However, this results in a wider filter bandwidth and a greater susceptibility to noise and cochannel interference. A more complete solution is to reduce the number of stages to achieve the correct sampling rate and then follow this by a one stage non-decimating filter designed to reduce the bandwidth to the desired value. The computational penalty of adding a single stage filter is not high since the filter is working on decimated data, not the full input.

When implemented this way, the appropriate halfband filter size is still determined using (65) except with the modification

$$M_2 = \frac{M-1}{D_{max}-1} + 1 \quad (68)$$

where D_{max} represents the maximum decimation rate (at any higher decimation rate the signal will be undersampled) and is given by

$$D_{max} = \frac{F_s}{F_{bw}} \quad (69)$$

The number of computations, C_1 due to the multistage halfband filter alone is given by (64), which is repeated and modified here as

$$\begin{aligned} C_1 &\approx \frac{7nN}{2} + nN(M_2 + 3) \left(1 - \frac{1}{D}\right) \\ &\approx \frac{7nN}{2} + nN \left(\frac{M-1}{D_{max}-1} + 4 \right) \left(1 - \frac{1}{D}\right) \end{aligned} \quad (70)$$

where D is the actual decimation rate ($D = 2^K$ where K is the number of stages).

In assessing the computational penalty of the final stage filter, both the number of input data values and filter size are required. The number of input data values feeding the final stage filter will equal the number of data values output by the multistage halfband filter and is found using (58) where $k = K$. Taking into account the latest relationship for M_2 expressed in (68), then (58) can be rewritten as

$$L_K = 1 + \frac{N-1}{D} - \frac{M-1}{D} \left(\frac{D-1}{D_{max}-1} \right) \quad (71)$$

which simplifies to

$$L_K \approx \frac{N}{D} \quad (72)$$

for large N . The filter size, M_1 , can also be derived from (71) since $L_K = M_1$ when $N = M$. Accordingly,

$$M_1 = 1 + \frac{M-1}{D} \left(\frac{D_{max}-D}{D_{max}-1} \right) \quad (73)$$

Using the derivation of (47) as a guide, and given that for large N the number of input data values L_K will approximately equal the number of output values (since it is a non-decimating filter), then the number of computations required for the final stage is given by

$$\begin{aligned} C_2 &\approx 2nL_K M_1 \\ &\approx \frac{2nN}{D} + \frac{2nN(M-1)}{D^2} \left(\frac{D_{max}-D}{D_{max}-1} \right) \end{aligned} \quad (74)$$

The total number of computations of the composite filter will be the sum of the contributions from each of the filters, or

$$C \approx C_1 + C_2 \quad (75)$$

Combining (70) and (74), and then simplifying yields the final result

$$\boxed{C \approx \frac{15nN}{2} - \frac{2nN}{D} + \frac{nN(M-1)(D^2 - 3D + 2D_{max})}{D^2(D_{max}-1)}} \quad (76)$$

4. Comparative Results

In the following two sections, the frequency response and computational requirements of the single stage, multiple stage (halfband and composite), and the spectral filters are compared and analyzed. These comparisons are carried out with emphasis on the demodulation of signals which are narrowband relative to the input data bandwidth.

4.1 Frequency Response

The filter frequency response, of an ideal filter would pass all signal energy within a desired passband, while completely rejecting all interfering signals and noise outside this band. Attempting to achieve this kind of performance comes at the cost of large filter lengths as illustrated in Figure 3. In this example, the frequency response of a polyphase filter⁴ is shown for different filter length. To make the results as general as possible, the filter length is defined as

$$M = \beta \left(\frac{F_s}{F_{bw}} \right) \quad (77)$$

where β is the filter length factor.

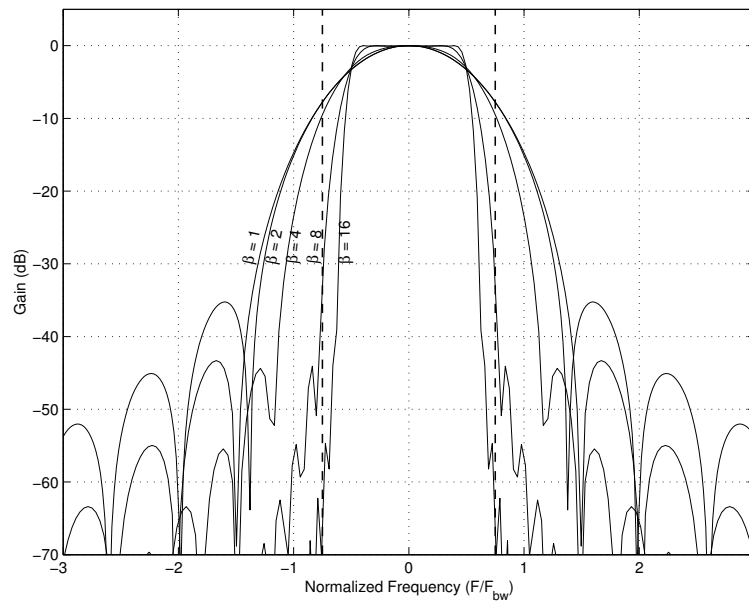


Figure 3: Filter response as a function of filter length.

The relevant features of the filter response are the flatness of the passband (the frequency region between -0.5 and 0.5), the roll-off rate in the transition band (i.e. the rapid drop in filter gain immediately outside), and the height of the sidelobes in the stopband. A flat response across the passband minimizes distortion to the desired signal

⁴The filter was designed based on the procedure outlined in [1] and using a Hanning window.

while a rapid roll-off and low sidelobes minimizes cochannel interference and noise. Better response can be achieved by increasing the filter length as illustrated in Figure 3. However, increasing filter length increases the processing requirements (i.e. C increases as a function of M in (40),(52),and (66)), hence the smallest filter length is usually chosen where acceptable performance is still achieved.

What is meant by “acceptable performance” will depend on the application as well as the signal environment. Specifying a maximum frequency width to the transition bands while maintaining reasonable passband flatness and stopband attenuation is the criteria often used. Given the trend observed in Figure 3, an approximately equivalent criteria requires a given attenuation at a frequency offset of $\pm 0.75 \times F_{bw}$ relative to the filter center frequency. The critical points are marked by the vertical dotted lines in Figure 3. In Figure 4, the filter gain at ± 0.75 is plotted for various values of β . From this figure, and choosing 30 dB criteria as a reasonable attenuation, acceptable performance is achieved when $\beta = 8$.

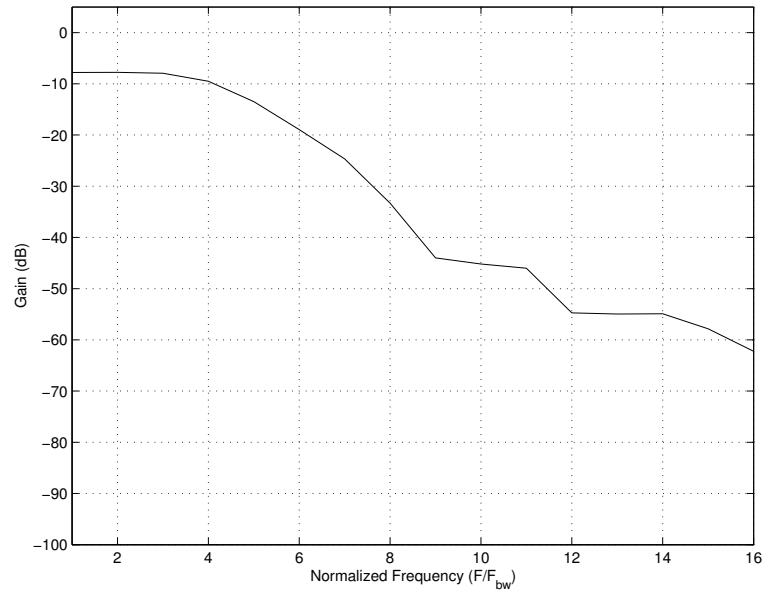


Figure 4: Filter gain at a normalized frequency of ± 0.75 as a function of filter length.

The choice of $\beta = 8$ is not meant to be considered as “the” choice since various engineering considerations come into play in its determination. For example, the choice of criteria parameters is as much based on the signal environment as on user preferences. A larger transition band requirement, for example, would lead to a smaller filter size. Additionally, there are some filter designs which generate lower sidelobes (hence better stopband attenuation) which might lead to a smaller filter size. The price, however, is poorer passband response and a larger transition band. The filter design chosen here was considered to be a reasonable compromise.

Using $\beta = 8$, and given $F_s/F_{bw} = 256$, examples of the frequency response of a single stage filter, a multistage halfband filter (a composite filter was not required for this

example), and the spectral filter are shown in Figure 5. The best performance is achieved using the single stage filter. This is not surprising since it has the greatest flexibility in design given that all M coefficients of the filter may be chosen independently. The least flexibility is found using the multistage halfband filter since only M_2 filter coefficients are available. By comparison, the spectral filter has $M/2$ weighting coefficients which can be independently modified if the constraint in (24) is followed. Performance-wise, the composite filter has a similar roll-off rate to the single stage filter but higher sidelobes. The spectral filter has a better response with lower sidelobes but a slower roll-off than the single stage filter.

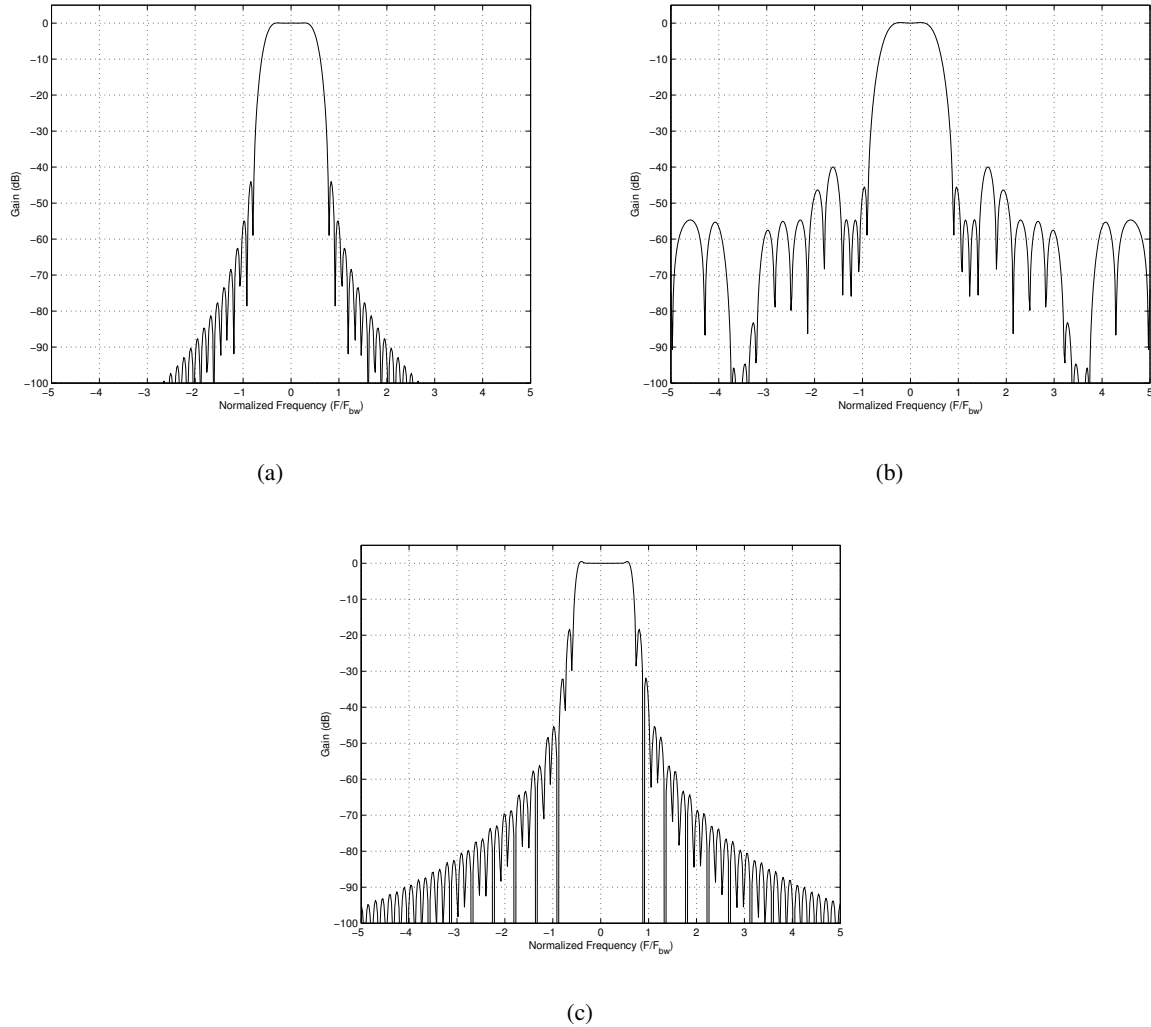


Figure 5: Frequency response of (a) the single stage filter, (b) the composite filter, and the (c) time/frequency spectral filter. The filter size factor was $\beta = 8$.

4.2 Computational Comparisons

Before comparing the computational advantages and disadvantages of the filters discussed so far, it is useful to define the relationship between the decimation rate and the signal bandwidth, which is,

$$D = \frac{1}{\alpha} \left(\frac{F_s}{F_{bw}} \right) \quad (78)$$

where α is the oversampling factor. For digital demodulation techniques, the complex sampling rate of the filtered signal will obviously have to be at least equal to the signal bandwidth ($\alpha = 1$). However, better demodulation performance can often be achieved using higher sampling rates (typically $\alpha \geq 2$).

Using $\beta = 8$ and $\alpha = 2$ as reasonable starting values, Figures 6 and 7 plot the effect of varying these values on the number of computations for the three filter types considered. The number of computations has been shown normalized with respect to the number of input samples. Additionally, it is also assumed that only one signal requires demodulation. The stepped nature of the curve for the multi-stage halfband filter is due to the fact that the halfband filter sizes are restricted to $M_2 = 5, 9, 13, 17, \dots$. Hence M_2 was selected by first using (68) and then selecting the closest legal value.

Examining the results, the multistage filter requires only about 25% of the number of computations of the signal stage filter in both Figures 6 and 7. The reasons will be discussed shortly.

Comparing the single stage filter to the spectral filter, the single stage filter required a smaller number of computations for lower values of β and α , but a greater number for higher values. This is due to the fact that increasing either the filter size or oversampling factor results in a linear increase in the number of computations for the single and multistage filters, but significantly less for the spectral filter. In fact, the main computational load for the spectral filter is generating the time/frequency spectral data (the first term in (40)) so that increasing the filter size results in a logarithmic increase in the computations (i.e. $\log M$) and increasing the oversampling factor (or decreasing the decimation rate) has only a minor effect.

Comparing the multistage filter to the spectral filter, the multistage filter required fewer computations in all cases shown. As with the single stage filter, there is a linear increase in the number of computations as a function of filter size. However in Figure 6, the filter size was not increased high enough for the spectral filter to obtain better results.

The computation equations representing the filters under evaluation (i.e. equations ((40), (52), and (76))), can be approximated by simpler forms which are more suitable for evaluation if certain conditions are met. Particularly if the desired filter bandwidth is a small percentage of the input sampling rate, and the oversampling factor is not large so that

$$D^2 \gg D_{max} \gg 1 \quad (79)$$

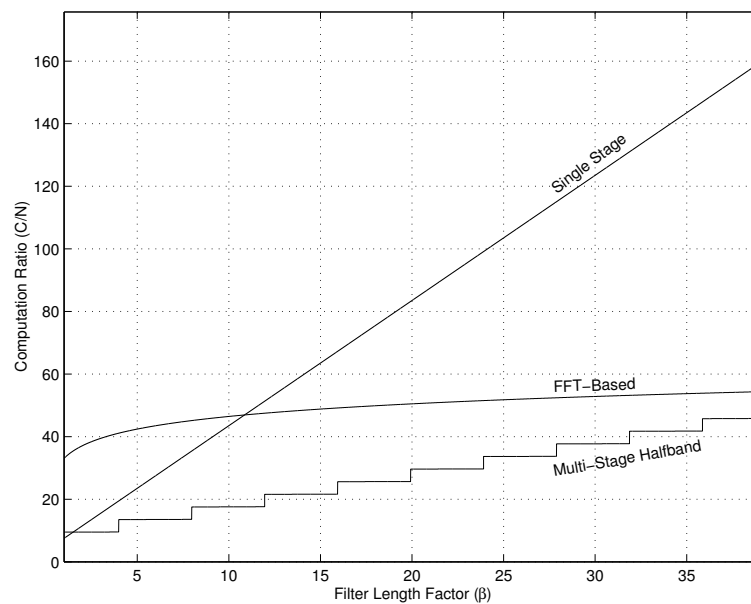


Figure 6: Computation ratio C/N as a function of the filter length for an oversampling factor $\alpha = 2$ and a single signal.

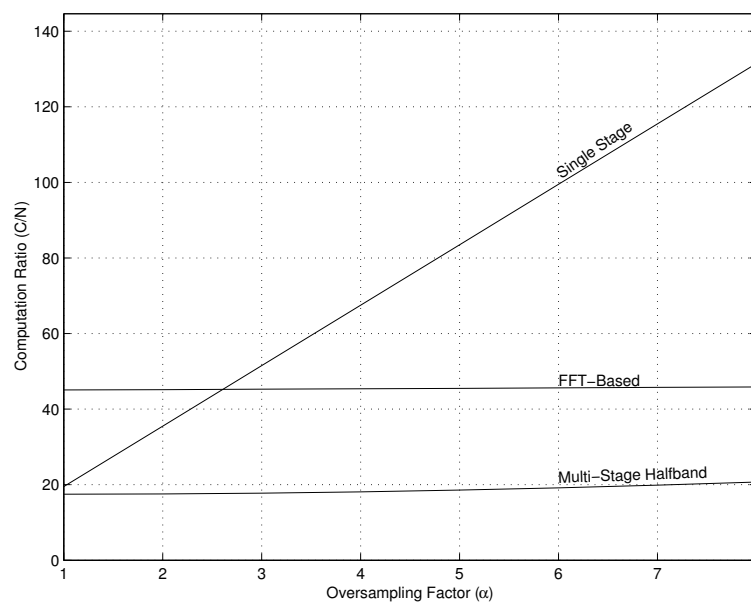


Figure 7: Computation ratio C/N as a function of the oversampling factor for a filter length factor $\beta = 8$ and a single signal.

Additionally, if the number of signals to be filtered is also not large ($D \gg n$), then the simplified computation rates are given by

$$\frac{C_{tf}}{N} \approx 4 \log_2 M \quad (80)$$

$$\frac{C_{ss}}{N} \approx \frac{7n}{2} + \frac{2nM}{D} \quad (81)$$

$$\frac{C_{cf}}{N} \approx \frac{15n}{2} + \frac{nM}{D_{max}} \quad (82)$$

which represent the spectral, single stage, and composite filters, respectively. For wideband systems processing narrowband signals, these assumptions are reasonable.

From these expressions, it is apparent that the lower computation rate of the composite filter compared to the single stage filter is primarily due to the advantage of zero coefficients in the halfband filter (note the extra factor of 2 in the second term of (81) compared to the second term of (80)) and the beneficial effect of doing the decimation as early in the filtering process as possible (another D/D_{max} improvement).

It is also clear from these expressions that (as expected) increasing the number of signals, n , results in a linear increase in the number of computations for the single stage and composite filters, but no significant increase in the number of computations for the spectral filter. Re-examining Figures 6 and 7 in this context, then for two or more signals the spectral filter will generally be faster than the single stage filter (e.g. for two signals the ratio C/N in either figure will be doubled for all single stage results making them greater, for the most part, than the corresponding FFT-based results which remain unchanged).

Comparing the spectral filter to the composite filter, with respect to the number of signals, is slightly more complicated. Mathematically, the spectral filter is faster when

$$\begin{aligned} C_{tf} &< C_{cf} \\ 4 \log_2 M &< \frac{15n}{2} + \frac{nM}{D_{max}} \end{aligned} \quad (83)$$

Given (69) and (77), then the condition can be rewritten as

$$\log_2 M < \frac{n}{4} \left(\frac{15}{2} + \beta \right) \quad (84)$$

Figure 8 illustrates this condition graphically over a range of different values for M , β , and n (note that in generating these curves, no attempt was made to ensure that M_2 was a legal value). Although these are approximate results, they suggest that for reasonable filter sizes ($M < 65536$) and passband response ($\beta \geq 8$), the spectral filter is faster when four or more signals are processed simultaneously.

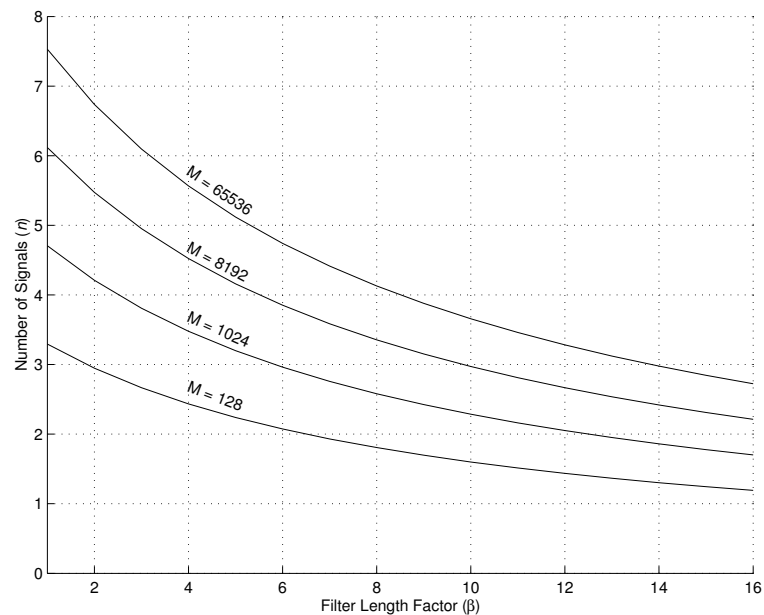


Figure 8: Parameter curves (β versus n) which mark the boundaries between faster processing for the spectral filter (above the curve) and faster processing for the composite filter (below the curve). Different curves are plotted for different filter sizes (M).

5. Conclusions

In this report a fast filtering approach suitable for applications involving the simultaneous demodulation of a number of narrowband signals captured in wideband digital data was introduced. This filtering approach is tailored for spectrum surveillance systems which first generate a time/frequency spectrum from the wideband data for the purposes of signal detection and tracking purposes. Using this spectral data, the data is then processed to produce the desired filtered and decimated signals at baseband frequencies. This is different from the normal approach where, for example, an FIR filter is used to process the input wideband data to produce the desired processed signal. In this report, the new approach is called the spectral filter.

The main advantage of the spectral filter is that the bulk of the processing involves the generation of the time/frequency spectrum which can be done efficiently using a fast Fourier transform (FFT). Subsequent processing to extract and filter the desired signals requires very little additional processing, hence there is relatively little difference whether one or a number of signals are extracted. By way of comparison, for normal time domain based filters the processing increases as a linear function of the number of signals.

The main disadvantage is that each time/frequency spectrum needs to be generated from blocks of data which overlap the previous data block by 50%. Although this also has advantages for signal detection, for the sake of speed, current systems such as

MiDAS use no overlap so that spectral filtering requires that the processing be doubled. For the time domain filters, there is no such doubling since the spectral data is ignored.

Two other filters were also introduced for comparative purpose, namely, a single stage FIR filter and a multistage FIR filter. The single stage filter has the greatest flexibility in terms of designing the filter frequency response. The multistage filter has considerably less flexible, but is computationally more efficient.

In terms of filter frequency response, the spectral filter falls in between the single stage and multistage filters. It is more flexible and has a better response than the multistage filter, but is still relatively inflexible when compared to the single stage filter.

In terms of computational speed, both the single stage and multistage filters are faster when only one signal is filtered. However when two signals are processed simultaneously, the spectral filter is faster than the single stage filter, and when four signals are processed simultaneously, the spectral filter is also faster than the multistage filter.

The assessment of comparative computational speeds is based on the assumption that implementing the spectral filter requires a doubling of the spectral processing already being carried out. If no spectral processing was previously carried out, then the relative performance will be worse, and the number of signals for which the spectral filter is faster becomes four and eight (instead of two and four stated in the previous paragraph). On the other hand, if the full spectral processing is already being carried out (i.e. the input data blocks overlap by 50%), then the spectral filter will be faster for any number of signals.

The assessment of comparative computational speeds is also based on ignoring various non-arithmetic processor activities such as memory transfers, pipelining, etc.. These activities could alter the relative performances of the various filtering approaches analyzed. In addition some filter approaches lend themselves better to efficient hardware designs which could significantly alter the relative performance results from those predicted by a simple number of computations assessment. Actual testing of the various filtering approaches on real systems would be required to investigate this.

The main conclusion is that for typical applications where the demodulation of four or more narrowband signals is desirable, the spectral filter is the preferred choice.

References

1. Read, W.J.L., (2002). Detection of frequency hopping signals in digital wideband data. (DRDC Ottawa TR 2002-162). Defence R&D Canada – Ottawa..
2. Stremmer, Ferrel G., (1977). *Introduction to Communications Systems*. Addison-Wesley Publishing Company, Inc., Don Mills, Ontario, Canada, 1977.

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.) <p style="text-align: center;">DEFENCE R&D CANADA – OTTAWA DEPARTMENT OF NATIONAL DEFENCE OTTAWA, ONTARIO, K1A 0Z4</p>		2. SECURITY CLASSIFICATION (overall security classification of the document, including special warning terms if applicable) <p style="text-align: center;">UNCLASSIFIED</p>
3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C or U) in parentheses after the title.) <p style="text-align: center;">A TIME/FREQUENCY SPECTRAL BASED APPROACH TO FILTERING (U)</p>		
4. AUTHORS (Last name, first name, middle initial) <p style="text-align: center;">READ, WILLIAM, J.L.</p>		
5. DATE OF PUBLICATION (month and year of publication of document) <p style="text-align: center;">DECEMBER 2005</p>	6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.) <p style="text-align: center;">29</p>	6b. NO. OF REFS (total cited in document) <p style="text-align: center;">2</p>
7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.) <p style="text-align: center;">DRDC OTTAWA TECHNICAL MEMORANDUM</p>		
8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.) <p style="text-align: center;">DEFENCE R&D CANADA – OTTAWA DEPARTMENT OF NATIONAL DEFENCE OTTAWA, ONTARIO, K1A 0Z4</p>		
9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant) <p style="text-align: center;">12PG11</p>	9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)	
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.) <p style="text-align: center;">DRDC Ottawa TM 2005-198</p>	10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)	
11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification) <p>(x) Unlimited distribution () Distribution limited to defence departments and defence contractors; further distribution only as approved () Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved () Distribution limited to government departments and agencies; further distribution only as approved () Distribution limited to defence departments; further distribution only as approved () Other (please specify):</p>		
12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.) <p style="text-align: center;">UNLIMITED</p>		

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

This report describes a new filtering approach suitable for demodulating narrowband signals captured in digital wideband data. The new approach is based on generating the filtered and downconverted baseband signals using time/frequency spectral data, rather than the input wideband data. For communications electronic support measures systems, this spectral data will normally be generated to facilitate signal detection and parameter estimation (i.e. frequency, bandwidth, bearing, etc.). Hence the new approach is able to take good advantage of this to reduce the overall number of computations. Through comparisons with more traditional FIR filtering techniques, it is shown that the new spectral filtering approach is computationally more efficient when multiple signals are to be simultaneously demodulated..

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

DIGITAL FILTER
FIR
FAST FOURIER TRANSFORM
FFT
MULTISTAGE
MULTIRATE
HALFBAND
WINDOW
HANNING
HAMMING
FREQUENCY SPECTRUM